CONNEXIONS™

The Interoperability Report

February 1988 Volume 2, No. 2

ConneXions - The Interoperability Report tracks current and emerging standards and technologies within the computer and communications industry.

In this issue:

ConneXions is published by Advanced Computing Environments, 480 San Antonio Road, Suite 100, Mountain View, CA 94040, USA. Phone: 415-941-3399.

ISSN 0894-5926

From the Editor

A European friend told me the other day that we should use the word *transition* rather than *migration* when we talk about changing protocol suites, since the latter word implies that we will be coming back (like the birds at the end of winter). As we prepare to *transition* from the DoD protocol suite to an ISO/OSI suite, it is interesting to look at some of the component protocols and compare their functionality. This month Rob Hagens of the University of Wisconsin compares the DoD IP with the corresponding ISO CLNP.

Many factors will determine the success of OSI, both in the US and worldwide. Michael Howard of Infonetics, Inc. takes the market analysis view and poses the question "Will OSI ever become a reality?"

There has been much talk about NetBIOS over TCP/IP. We asked John Romkey of FTP Software to give us another "inside look", answering the question "Just what is NetBIOS, anyway?"

Internetworks are growing every day. The announcement made recently by the National Science Foundation about its agreement to link its network with NASA's marks an important step in the effort to develop an interagency internet. See page 14 for more details.

Protocol testing continues to be an area of importance and several organizations are preparing testing labs which will exercise protocols at layer 4 and below. We will be bringing you more information about these efforts as they emerge. Meanwhile, the industry is rallying behind Sun's Open Network Computing architecture at layers 5, 6 and 7. A major connectivity test, known as the *Connectathon*, recently took place with participation from over 50 vendors. We bring you a report on page 15.

It is time once again to remind you that RFC documents are available from the Network Information Center (NIC) at SRI International in Menlo Park, California. Online copies may be obtained via anonymous FTP from the RFC: directory on host SRI-NIC.ARPA (10.0.0.51 or 26.0.0.73). Filenames are of the form RFC:RFCnnnn.TXT, where "nnnn" is the RFC number. Hardcopies may be ordered by calling the NIC at 800-235-3155 or 415-859-3695. CCITT and ISO documents are available from Omnicom, Inc. in Vienna, Virginia. Their phone number is 703-281-1135.

A Comparison of the DoD IP and ISO CLNP

by Rob Hagens, University of Wisconsin

Introduction

Connectionless network layer protocols provide the means to transmit a self-contained data unit ("*packet*") without establishing, maintaining or releasing a connection. These packets may be sent directly to a destination, or they may be relayed through intermediate gateways towards a destination. Although not implied by the word "connectionless", many of these protocols allow packets to be discarded during transit; the associated source does not guarantee delivery.

One well known example of this type of protocol is the Department of Defense Internet Protocol (described in RFC 791, or MIL-STD-1777). Known as IP, this protocol has been employed successfully by the Academic/Research Internet for a number of years.

Recently, a new protocol has appeared: the International Organization for Standardization (ISO) Connectionless Network Protocol, in this article referred to as CLNP. CLNP is described in ISO 8473.

In this article, these two protocols are compared. Although they provide roughly equivalent functionality at the service interface level, they diverge at the protocol level. For example, they differ in the way that a packet is structured; IP addresses are fixed length whereas CLNP addresses are variable length.

General

The formats of an IP packet (Figure 1) and a CLNP packet (Figure 2) are very similar. Each packet consists of a header followed by data. Although both packets may contain the same amount of data (65535 bytes), the maximum size of a IP header (60 bytes) is much smaller than the maximum size of a CLNP header (255 bytes).

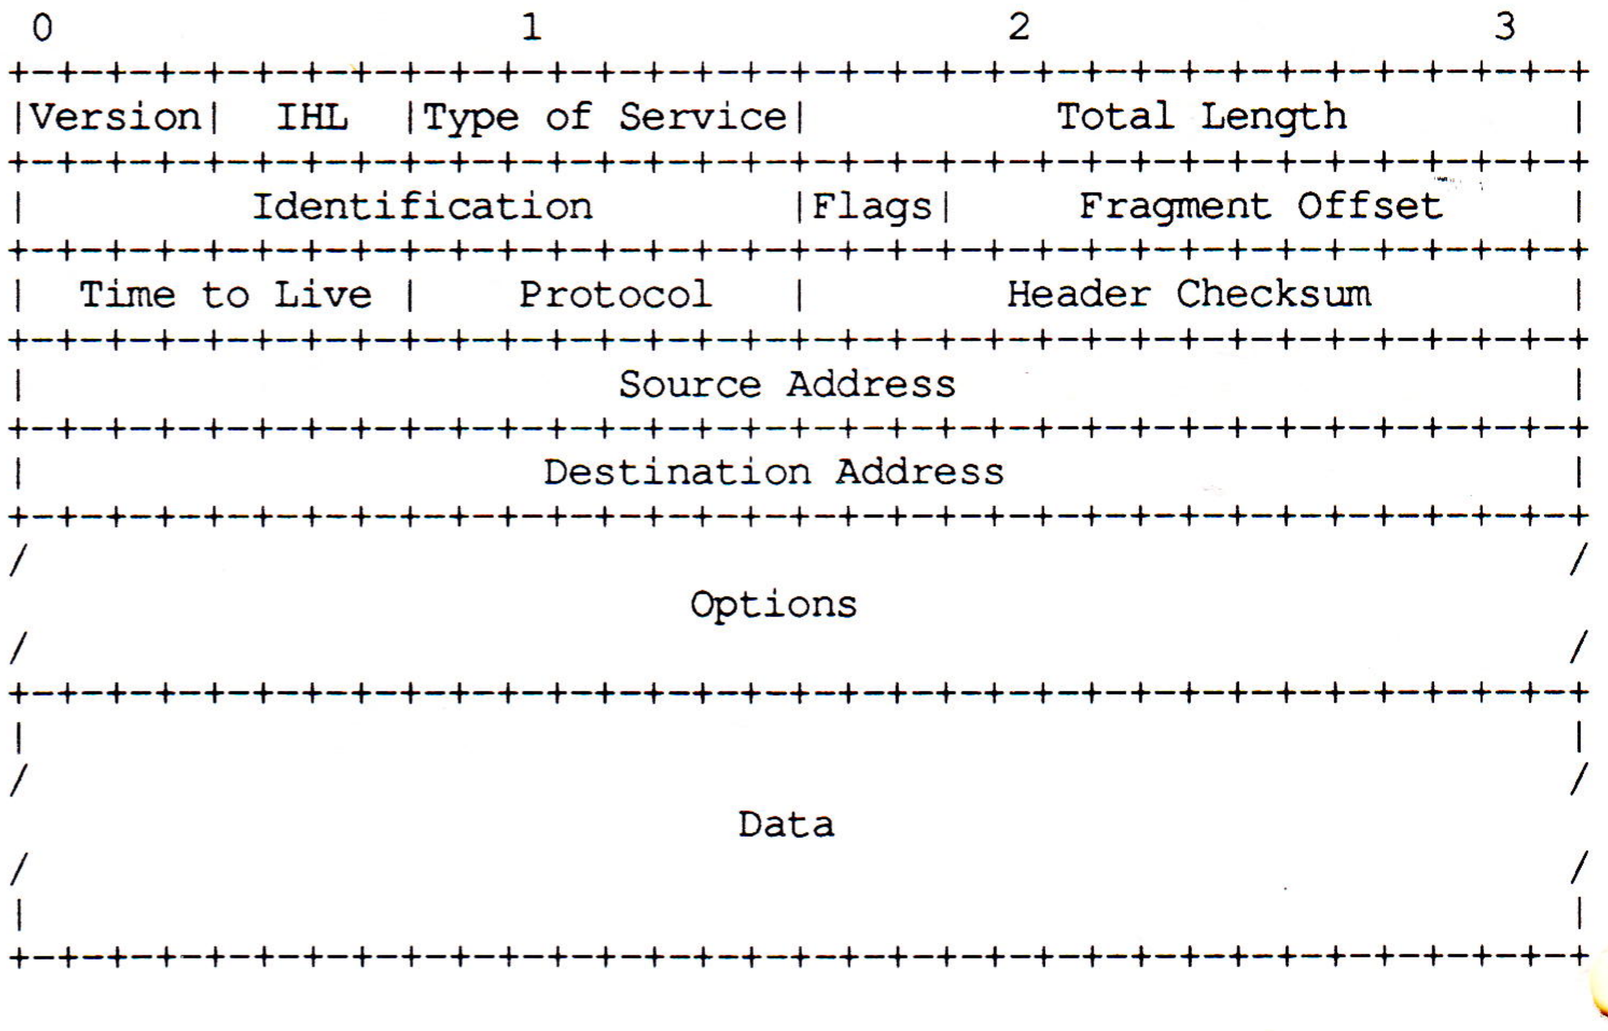

Figure 1: (Courtesy of RFC 791).
The format of an IP packet Internet Datagram

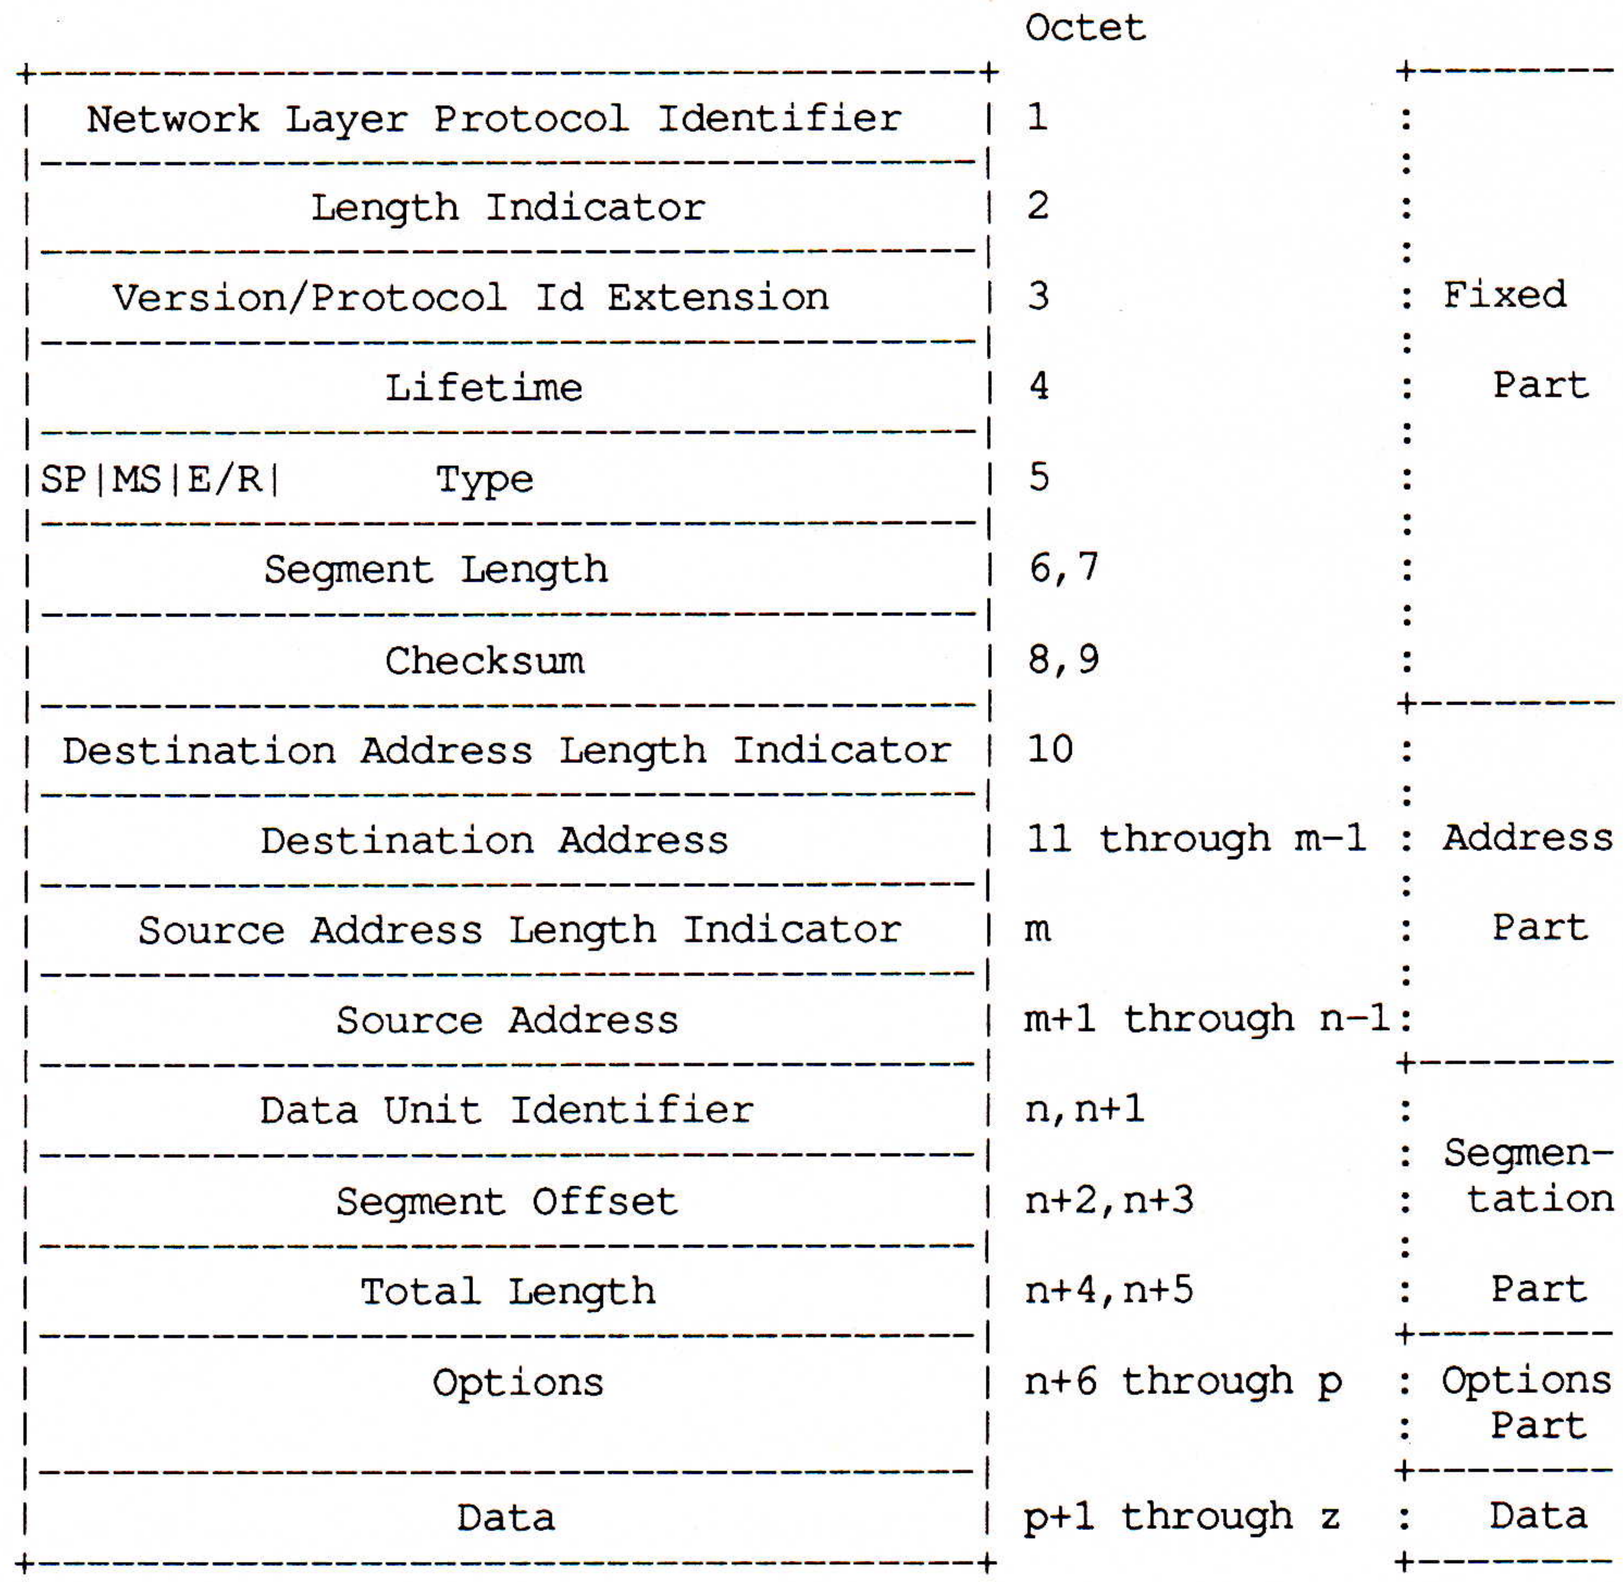

Figure 2: (Courtesy of RFC 986). The format of a CLNP packet.

Checksum

In order to detect errors introduced into the packet by transmission over an unreliable network, both protocols include a *checksum* computed on the packet header. IP uses the TCP checksum (one's complement of the one's complement sum) whereas CLNP uses the stronger, but more costly TP checksum (Fletcher's Algorithm). Optionally, CLNP checksum validation by the receiver can be turned off by specifying a value of 0 in the checksum field when the packet is created.

Header format

In both protocols, the header contains a fixed part followed by a variable part. Roughly speaking, the fixed part contains information every packet must carry, while the variable part contains optional parameters not necessarily present in every packet. One of the major differences between IP and CLNP is that the amount of information fixed in an IP packet is far greater than the amount fixed in a CLNP packet. For example, addressing and fragmentation information appears in the fixed part of an IP header; this information resides in the variable part of a CLNP packet header.

Both protocols contain a Time-To-Live (or "Lifetime", in ISO terminology) field in the header. This field applies an upper bound on the time a packet may exist in the network. The granularity of this field is seconds for IP and half-seconds for CLNP.

continued on next page

Comparison of DoD IP and ISO CLNP *(continued)*

A controversial difference between the IP header and the CLNP header is that the CLNP header does not contain an upper layer protocol identifier. This identifier indicates which higher layer protocol should receive the contents of the packet. This identifier is absent from CLNP because it is believed that higher layer identification is an addressing issue, not a protocol issue. The common solution to this problem is to include a "transport selector" field in the destination address.

Addresses

The differences in addressing between IP and CLNP occur in two areas: address format and the location of the address within the header. IP packets use the 32 bit Internet address format located in the fixed part of the packet header. CLNP packets use the ISO NSAP (Network Service Access Point) address format as specified by ISO 8348 (Network Service Definition), Addendum 2. Since these addresses are variable in length, they appear after the fixed part of the CLNP header. The extent of each address is indicated by a preceding length byte.

Fragmentation and reassembly

Both protocols support the concept of fragmentation (or "segmentation", in ISO terminology). Fragmentation allows a large packet to be broken into smaller pieces when it is too large to be transmitted by the network interface. Both IP and CLNP support the notion of a "don't fragment" flag which can be set on a per-packet basis. When this flag is set, fragmentation is not allowed. Since the fragmentation information does not reside in the fixed part of the CLNP header, this information may be omitted (saving 6 bytes) when the "don't fragment" flag is set. In ISO terminology, omitting the segmentation information is known as using the "non-segmenting protocol subset".

Reassembly is the process by which packet fragments are merged together into their original form. Since these protocols do not provide any delivery guarantee, it is possible for packet fragments to be lost. In this case, the remaining fragments awaiting reassembly must be discarded. Both protocols use a timer to determine when to discard the fragments awaiting reassembly.

Optional parameters

Optional packet parameters may be used to convey information that is not necessarily transmitted with every packet. These parameters are set at the discretion of the creator of the packet. In both protocols, these parameters are transmitted within the variable part of the header.

The format of the optional parameters -- parameter identifier, parameter length, parameter value -- is similar in both protocols. Thus, when a packet containing options is received by either protocol, the optional parameters must be parsed. It is important to note that parsing an optional parameter is more costly than accessing a field within the fixed part of the packet header.

Although the decision to send optional parameters is up to the creator of a packet, IP mandates that all implementations *shall* support every optional parameter. This is in contrast to CLNP, where the support of certain optional parameters is itself, optional.

A second difference in the handling of options can be seen when one considers the fragmentation of packets containing optional parameters. Each option in an IP packet has associated with it a bit that determines whether the option is copied when the packet is fragmented. In CLNP, all options are copied with every packet.

The next sections will describe differences between the specific options supported by IP and CLNP. In general, both protocols support the same type of options with the exception that there is no timestamp option in CLNP.

Security option

Both protocols support the notion of a *security parameter*. It is used to associate a packet with a particular security level. In IP, the security parameter has a fixed length of 11 bytes. CLNP employs a variable length security option which may be one of three types: source address specific, destination address specific, or globally unique.

Source Route option

Source routing can be used in order to predefine the route that a packet will take through a network. In both protocols, the source route takes the same general form: a list of addresses representing hosts which must be visited in order. Both IP and CLNP support the concept of *strict* ("complete", in ISO terminology) and *loose* ("partial", in ISO terminology) source routing. Strict source routing forces a packet to visit *only* hosts listed in the source route while en route to the destination. Furthermore, those hosts *must be visited in the order listed*. Loose source routing is like strict source routing except that the packet may visit any number of hosts while en route to the next host in the source route list.

A difference between the IP source route and the CLNP source route is that the IP destination address field is replaced by the next host listed in the source route. This in contrast to CLNP, where the original destination address remains unchanged in the packet header.

Record Route option

Both protocols have a *Record Route* option. This parameter permits hosts to log their addresses in the packet as it travels towards the destination. However, the protocols differ in the order in which the route is created. IP specifies that the addresses are recorded in traversed order. The resulting list must be reversed in order to be used as a source route from destination to source. CLNP, on the other hand, records route in reverse order. Thus a CLNP recorded route may be used by the destination as a source route back to the source. A final difference between the record route option in IP and CLNP is that IP provides an implicit record route whenever a source route is used: when a source route address is removed from the list, it is replaced with the address of the host just visited.

Quality of Service option

One of the most striking differences between IP Type Of Service information (TOS) and CLNP Quality Of Service (QOS) information is the location of the parameter. In IP, the parameter is located in the fixed part of the packet header. In CLNP, the parameter is located in the variable length, options part. IP TOS information (fixed in length) is defined in terms of precedence, delay, throughput, and reliability. In contrast, CLNP QOS (variable in length) is defined in terms of source address specific information (semantics undefined), destination address specific information (semantics undefined), and globally unique information.

continued on next page

Comparison DoD IP and ISO CLNP *(continued)*

Since CLNP quality of service information is variable in length, users of CLNP have far more flexibility in defining quality of service semantics. However, this flexibility has a price; the information must appear in the variable part of the header, thereby making it more costly to locate.

One of the bits in the globally unique information that has no counterpart in IP is the *"Congestion Experienced"* bit. The purpose of this bit is to indicate that the packet experienced congestion on its way toward the destination. Ideally, a transport entity that received a sufficient number of packets with this bit set would take steps to reduce its network traffic.

Finally, there is no counterpart in the CLNP Quality Of Service parameter to the precedence portion of the IP Type Of Service parameter. Rather, CLNP defines a distinct option called the Priority Option.

Errors

The networks over which these protocols operate may not be error free. Thus, each of these protocols employs a mechanism to flag errors encountered. IP reports errors via a separate protocol, the Internet Control Message Protocol (ICMP, defined in RFC 792) whereas CLNP reports errors via an internal Error Report function. The Error Report function of CLNP utilizes a packet distinct from the normal data packet.

The ICMP and CLNP error report function overlap in the following areas: they both have a means to report that a packet has been discarded because (1) a destination is unreachable; (2) the Time-To-Live (or Lifetime) field has reached zero; or (3) a parameter problem (or "syntax error", in ISO terminology) has been detected while parsing the packet.

The remaining ICMP functions have no equivalent in CLNP. For example, the ICMP redirect function has no counterpart in the CLNP error report packet. Instead, the redirect function is handled by the ISO ES-IS routing protocol (described in ISO DP 9542). The ICMP Source Quench, Echo, Timestamp, and Information Request functions have no direct counterpart in the ISO network layer.

Summary

It is clear from the similarity of features that the designers of CLNP used IP as their model. However, when building upon this model, the designers opted for *variable length*, flexible parameters in order to satisfy the diverse needs of an OSI environment.

ROBERT A. HAGENS received an M.S. in Computer Science from the University of Wisconsin-Madison in 1984. Hagens is currently an Associate Researcher in Network Communications at the University of Wisconsin-Madison. His current work involves implementing the ISO Session and Network layer protocols and UNIX kernel systems support for the Wisconsin ISO project (ISO protocol implementation for the IBM RT/PC workstation in a Berkeley UNIX environment). He is an observer in ANSI Task Group X3S3.3. Previously, Hagens worked on the WISCNET project (DoD protocol suite for IBM VM CMS systems) including implementation and performance improvements to FTP, UDP, and TCP/IP.

Just what is NetBIOS, anyway?

by John Romkey, FTP Software

Recently, IBM PC "NetBIOS" has acheived some prominence in the TCP/IP community. This article discusses just what NetBIOS is, why anyone would want to use it, and how it can work with TCP/IP.

A brief history of NetBIOS

NetBIOS was originally created by Sytek as an interface to their intelligent broadband network interface. The IBM PC normally has BIOS (the Basic I/O System) in ROM. BIOS is a set of drivers that provide simple hardware support for standard equipment on the PC (for instance, drivers for the keyboard, disk controller, printer and display). Sytek's NetBIOS is the equivalent of the IBM BIOS, but for their network interface.

After IBM started selling Sytek's broadband network as "the PC Network", they adopted NetBIOS as the convention for communicating with the network. Microsoft created "Microsoft Networks", which provides NFS-like filesystem sharing abilities, but tailored specifically for IBM PC's running PC/MS-DOS. IBM turned that into the "IBM PC Network Program" (for the Sytek interface) and "the IBM PC LAN Program" (for the IBM Token Ring). The IBM Network Programs bear a similar relationship to Microsoft Networks as IBM PC-DOS bears to Microsoft MS-DOS.

These programs all accessed the "network" via the conventions originally implemented in the Sytek NetBIOS, except that IBM added a new call - Find Name - to the Token Ring version.

A programming convention

NetBIOS is actually only a specification of how to talk to the network, not what you're actually talking to. The NetBIOS specification says that to perform a certain network operation, you load various processor registers with certain values and perform a software interrupt 5C (hex). The network layer handles this, performs the desired operation, and returns a status to you.

While the NetBIOS specification doesn't really force you to use any specific network layer, it does place fairly rigid constraints on what services this layer must provide. The network layer must provide unreliable (reception by the destination is not guaranteed) datagrams, and reliable sessions (connections). Datagrams have maximum lengths which vary by implementations, and they correspond nicely to UDP datagrams. Sessions correspond to TCP streams, except that session closes have slightly different semantics from TCP closes, and sessions have record marks in them whereas TCP streams are raw byte streams without record marks in them.

NetBIOS Names

The biggest difference between the NetBIOS view of the network and the TCP view of the network is the way names work. NetBIOS Names are 16 byte binary names. NetBIOS Names really name processes, not computers. For instance, a NetBIOS LAN might have a boot server on it carrying the name "IBMBOOTSERVER". When the boot server program runs on the PC, it calls NetBIOS and tells it that it wants to use the name "IBMBOOTSERVER" and its NetBIOS then communicates with the other NetBIOSes on the LAN, and verifies that nobody else is using that name.

continued on next page

Just what is NetBIOS, anyway? *(continued)*

When some other machine wishes to (in this case) boot, it will try to discover where that name is and send its packets to that computer. The user's application is spared having to worry about which computer is actually the boot server. It's sort of like having UDP and TCP port numbers without having IP addresses, with the major difference being that no two machines on the same LAN can have the same "port number" (NetBIOS Name).

What does this have to do with TCP/IP?

After IBM blessed NetBIOS by building more programs around it, more vendors developed applications on top of it (some database programs use it for record locking on shared databases and some terminal emulation programs work with it as well), and more vendors tried to supply it.

Sytek uses NetBIOS to interface to a private protocol described in the back of the IBM PC Network Technical Reference Manual (which, incidentally, is *the* place to read about NetBIOS calling conventions). IBM uses it above its own private protocol on the token ring. DEC has a version that uses DECnet to access the network; several vendors have implemented it over what exists of the ISO protocol suite; Novell and other proprietary PC network vendors have implemented NetBIOS over their protocols. So, naturally, there are a number of people working on it over TCP/IP as well.

NetBIOS sessions and datagrams map fairly easily onto TCP and UDP; a little bit of protocol must be interposed to carry some information that there is no other convenient way to communicate (such as NetBIOS Names and frame boundaries over sessions). The biggest problem is that NetBIOS Names don't act anything like TCP/IP addresses, so a substantial amount of protocol must be put in place here to make things work.

The semantics of NetBIOS Names make it difficult to use them in an internetwork. NetBIOS Names are very dynamic; they can move from host to host quickly; a computer using a name can vanish from the system with no notification. The easiest way to implement NetBIOS Names involves using broadcast, which is at best ineffective and at worst destructive over an internetwork (although it's reasonable on a LAN as long as not much traffic is broadcast).

Three kinds of nodes

RFC's 1001 and 1002 describe the standard way to use NetBIOS over TCP/IP. They provide for three forms of NetBIOS: "B-nodes", "P-nodes" and "M-nodes". B-nodes are broadcast-only (and restricted to a single LAN); P-nodes are point-to-point (and provide internetworking) and M-nodes are mixed broadcast and point-to-point. B-nodes implement the name protocol by broadcasting packets over the LAN. P-nodes use "NetBIOS Name Servers", and M-nodes use a mix of LAN broadcast and name servers.

Name protocol

The actual Name protocol in RFC 1001/1002 is an incompatible subset of the Domain Name Protocol. While there is a certain amount of technical elegance to using an existing design to solve a new problem, the Domain Name Protocol was, in this case, overkill. NetBIOS adds totally new operations to it, not using any of the old ones, thus all that NetBIOS gets out of the Domain Name is the packet format, which is much more complicated than is necessary.

Why use NetBIOS over TCP/IP?

Running NetBIOS over TCP/IP means that any system you can reach via TCP you can also reach via NetBIOS. (If that system supports the NetBIOS protocols from RFC's 1001 and 1002, and the applications protocols which use NetBIOS). You could expect to be able to access files from a UNIX system supporting NetBIOS and SMB (the protocol used by Microsoft Networks for filesystem sharing) while running FTP at the same time. Since the UNIX system will presumably already have TCP/IP in place, it's easier and more general just to bring up NetBIOS over TCP/IP than it is to have to implement whatever non-TCP protocol you choose to run NetBIOS on.

Beware that only P-node and M-node NetBIOS's can communicate with computers that aren't on the same LAN, and most implementations currently announced or available are B-node implementations.

The original implementations for TCP/IP were done by Excelan and Ungermann-Bass. Mitre stepped in as a mediating agency to help bring about a compromise on how NetBIOS would be implemented over TCP/IP, and RFC's 1001 and 1002 were born.

Where can you get NetBIOS?

The following companies are either currently shipping or have announced their intentions to develop an RFC 1001 and 1002 compliant NetBIOS for the IBM PC: Bridge Communications, Excelan Inc., FTP Software, Network Research Corp., Communications Machinery Corp., Syntax, and Ungermann-Bass. (This isn't meant to be an exhaustive list). Other vendors are also supplying NetBIOS on systems other than PC's (for instance, UNIX).

Several pieces of technical documentation exist. The NetBIOS specification is the IBM PC Network Technical Reference. RFC 1001 and RFC 1002, available from the NIC, describe how to implement NetBIOS over TCP/IP.

The author owes many thanks to Karl Auerbach for his help with NetBIOS problems in the past and for helping to clean up some points in this article.

JOHN ROMKEY received his B.S. in Computer Science from MIT in 1985. While there, he worked with Prof. Jerome Saltzer and Dr. David Clark for three and a half years on the PC/IP project, a popular public domain implementation of TCP/IP for IBM PC's. After graduation, he spent a year as a staff member with Dr. Clark and moved on to help found FTP Software, Inc., where he remains as Chief Architect. Perhaps sometime he'll get back to writing Science Fiction.

Will OSI ever become a reality?

by Michael Howard, Infonetics, Inc.

In the Fall of 1986, the industry was buzzing about the promise of the Open Systems Interconnect (OSI) standards to fulfill the market desire for interconnection of heterogeneous computing environments. It was reported that the federal Office of Management and Budget (OMB) was making plans to require OSI protocols in future government procurements. Many vendors were talking about future OSI products and how their products would be able to work with other vendors' computing systems. The problems and costs of multivendor computing within corporations and government agencies would dissipate by use of common applications and networking protocols. Where is the promise of OSI today? What will happen in the next two years? Will OSI achieve the promise of internetworking or will it ever become a reality?

OSI

OSI stands for the Open Systems Interconnection reference model developed by the ISO (International Standards Organization). However, the term is used in the industry to cover the much broader area of all the protocols and other ISO standards that are specified to operate under the OSI 7-layer model. Some of the most well-known OSI standards are the X.25 networking protocol and the X.400 electronic mail messaging standard. We do not intend to describe the OSI model, the OSI protocols and the various standards bodies here -- those are available from many sources. We do want to look at the genesis of the OSI standards from a market perspective, look at their status today, and look at what the future holds. But first, a little background information about internetworking.

Interconnection beginnings

In the beginning of interconnection, there was Arpanet. In 1966, the first host-to-host version of the government-funded Arpanet connected dissimilar computing environments. This early version gave way to the TCP/IP Arpanet in the early 1970s. The purpose of Arpanet was to connect the dissimilar computers at various universities, defense contractors and the defense agencies that were paying for it's development. About the same time in the mid 1970s, IBM announced the first version of SNA, or Systems Network Architecture, to tie together their disparate set of computing machinery. The same interconnection needs recognized by users and vendors brought about the establishment of ISO standardization efforts for Open Systems Interconnection in 1977.

Interconnection functions

What are the functions provided by interconnected networks? Interconnection functions can be organized into distribution and processing. Information can be distributed to dissimilar environments. Examples are electronic mail and file transfer. Processing occurs on a dissimilar computer either through terminal mode or through a host-to-host mode using program to program communication. The most typical application today is access to database information. In the future, we will see more specialized computing systems with high-speed database access or with artificial intelligence applications to process information from other database machines.

Why interconnect?

Networks can be connected by bridges or gateways. Bridges connect similar networks usually in order to extend the range (physical distance covered) of the network. Bridges can be thought of as switches or relay points within a similar network. The interesting connection, then, is a gateway that connects two dissimilar networks. If two completely different networks such as SNA and DECnet are connected via a gateway, then all 7 layers of each protocol stack must be implemented in the gateway. This is a complicated translation of addresses, data formats, user identification, internetwork routing, and other protocol considerations. Such gateways are inherently inefficient and give rise to the demand for an interconnection standard. If the bottom three layers of a protocol stack are the same, e.g., X.25, then such a gateway becomes much simpler.

Local Area Networks

Probably the most potent factor today in driving the demand for interconnection of networks is the proliferance of Local Area Networks (LANs). As LAN installations increase, the users want to connect to other LANs and departmental, mid-range and mainframe computers. The phenomenal proliferation of microcomputer workstations has reached such a critical mass. This critical mass represents a large investment and an accompanying corporate conern over maximizing productivity gains from that investment. As more workgroup tools are available to workers, and the information needs increase, LANs are being purchased at a growing rate. In fact, LANs are the fastest growing sector of the computer industry, with over 50% of LAN installations occurring in the last two years.

Interconnection standards

The world is divided in two: IBM and non-IBM. In the IBM world, there is SNA with some connections to the outside world. IBM has announced several interfaces to OSI protocols and applications; however, SNA will always be the backbone network service standard sold by IBM. IBM is one of the many vendors filling the committee positions of the standards organizations. OSI demands are strong in Europe and IBM still gains healthy revenues from the European market, so we believe that IBM will continue to strongly support OSI interfaces from SNA. There is movement to place the APPC LU 6.2 interface specification in the OSI standards. IBM has reportedly committed to put LU 6.2 in the public domain if adopted as part of OSI.

In the non-IBM world, there are TCP/IP and OSI as possible interconnection standards. And every major computer hardware and network vendor connects to SNA. XNS was a potential standard as recently as three years ago, but it has been outvoted in the marketplace and is dead. Some key LAN vendors originally implemented XNS, but have now switched over to TCP/IP. DEC has recently announced intention to implement an OSI stack. Other vendors are planning to implement OSI protocols alongside their existing TCP/IP stacks.

Factors pushing OSI

There are several possible factors discussed below that will push OSI into wider use. The desire for interconnection and internetworking is real among vendors and users, and this is the basis for OSI in the first place, but OSI is not the only answer. We will discuss those factors that would allow OSI to win out over other solutions.

continued on next page

Will OSI ever become a reality? *(continued)*

OSI is strong in Europe. The Europeans are usually ahead of the US in supporting standards, and some European governments are requiring OSI protocols and applications in major procurements.

X.400 Electronic Messaging

The demand for messaging may push the X.400 electronic mail standard into popular use. This alone does not give us networking interconnection in the lower four transport layers, but makes it more likely. Since X.400 is an application and presentation layer standard, it can be implemented on top of many other networks. Indeed, DEC, GTE, and other vendors have announced X.400 products. IBM has, of course, announced certain X.400 interface products.

A number of vendors have announced OSI protocol support in some form or other. At least two companies, Retix and Touch Communications, are dedicated to producing OSI protocol products. DEC is committed to offering OSI transport protocols.

Some parts of the marketplace may shy away from TCP/IP and support OSI because it is not produced by a standards body. It may not need a standards body approval, but it needs some kind of certification such as passing the DOD TCP/IP test suite.

GOSIP

Finally, we believe the strongest push for OSI has not occurred yet. Some market event is needed, and the most likely event is for the federal government to require OSI for networking procurements. A year ago, the OMB was thought to be close to issuing a requirement for each cabinet-level department to draw up a plan for conversion to use of OSI. This document has never materialized. The Government Open Systems Interconnect Profile (GOSIP) committee is supposed to issue a report on the government use of OSI products and protocols in a few months. This may or may not happen since the basis for the conversion to OSI rests on the assumption of an active and competitive OSI market. The industry is clearly not there yet. This event could still occur with 24 months.

Factors retarding the growth of OSI

There are several factors retarding the growth of OSI. Many of the standards are not complete. For example, the critical directory structure part of the X.400 will not be a specified standard until at least 1989, leaving each vendor to create a different structure in today's X.400 implementations. If it takes so much time to finalize the messaging standard, then it will take a long time to complete other communications protocols.

Cost

Finalization of standards is just the first step. Vendors must invest their development dollars and, maybe even more importantly, their development staff resources in the implementation of these standards. Early partial implementations of transport protocols have been lengthy and expensive, leading vendors to consider waiting until the market demand is clear.

Scope

Most of the standards are too all-encompassing and/or have no efficient bypass of unwanted function. Vendors will choose which parts to implement and which parts to leave until later. Current implementations of transport protocols are very slow. And different vendors' implementations on the same hardware have yielded widely differing performance results, all slower than tolerable.

In the IBM world, the Systems Applications Architecture (SAA) is the new IBM strategy for a broadened definition of heterogeneous connectivity of IBM's disparate computing environments. We believe that SAA will become a standard, and it will come much more quickly than did SNA. SAA is a communications standard and much more. It is a connectivity standard in that it defines user access and program portability. Yes, it defines a standard access at the communications level, but also defines a standard access at the program and user levels. This higher level approach to connectivity will mollify the IBM world's need and thus leave no requirement to consider other solutions like OSI. With that said, it appears that IBM still has a chance to be directly connected with an OSI standard if the APPC/LU6.2 interface is adopted.

TCP/IP works *now*

Finally, the most prominent factor is the need for interconnection now. The market must wait for OSI to be defined, to be implemented, to be made efficient, and then to be made competitive by many vendors. Meanwhile, TCP/IP has been in use and has been maturing for over a decade. There are a growing number of vendors introducing TCP/IP products. There is a large, growing installed base of TCP/IP users who, without a compelling reason to switch to OSI, will be saying, "don't fix what ain't broke".

Whither OSI?

What then about the future of OSI? We believe that there is a two-year window of opportunity for OSI to gain a strong foothold in the marketplace. Further, we believe that there must be some market event to drive the need for OSI interconnection products. As a given, IBM will continue to dominate the IBM marketplace and solidify it with its SAA. Additionally, IBM may make more inroads into the interconnection marketplace if LU6.2 becomes an OSI standard.

In the non-IBM world, TCP/IP is rapidly growing as an important interconnection standard, not only in its historical government, university and Unix markets, but in the commercial sector as well. In fact, Infonetics believes that the growth of the TCP/IP marketplace is on the beginning of an exploding growth curve. Infonetics estimates that the TCP/IP market will grow to a $831 million market by 1989.

The big question for OSI is: will the market-driving event occur within the next 24 months? If such a market event does not occur, then TCP/IP and IBM's SAA/SNA will dominate the interconnection marketplace.

MICHAEL HOWARD is executive vice president and senior analyst for internetworking at Infonetics, Inc. Infonetics is an international consulting company providing clients with publications, custom research and analysis services, and industry conferences, with offices located in Santa Clara, California.

NSF and NASA agree to link computer networks

The National Science Foundation (NSF) announced on January 4th, 1988 the signing of an agreement with the National Aeronautics and Space Administration (NASA) to share high speed communications lines, an effort that will link university researchers now connected to NSF's national computer communications network to data bases and supercomputers at NASA laboratories, saving hundreds of thousands of dollars that might otherwise be wasted in duplicated efforts by the two agencies.

Collaborative research

The agreement is in accord with a report just released by the White House Office of Science and Technology Policy (OSTP). The report, "A Research and Development Strategy for High Performance Computing," recommends improvements in networking to enhance U.S. leadership in the field and to provide the linkages needed for collaborative research by scientists working at different institutions.

"NSF and NASA have agreed to work together to identify ways to satisfy the technical requirements of networking through research, engineering, and implementation," said Steve Wolff, NSF Division Director for Networking and Communications Research and Infrastructure. "We are acting on the report's recommendations, and expect to indentify some cost effective solutions to networking problems as a result."

Planned connections

Three NASA facilities will be linked to existing NSF regional networks, which in turn are connected through a national backbone network. The Goddard Space Flight Center in Greenbelt, Maryland, will be linked to the Southeastern Universities Research Associates Net (SURANET); the Ames Research Center in Mountain View, California, will be linked to the Bay Area Regional Research Net (BARRNET); and the Johnson Space Flight Center, in Houston, will be linked to SESQUINET, a regional network in Texas.

Authorized scientists will be able to remotely access and use NASA data in their research, and can apply for time on NASA supercomputers. NASA-funded scientists at universities served by NSF regional networks will be able to communicate and collaborate with their colleagues at the NASA centers.

Interconnecting agencies

NSF already shares networking facilities with the Office of Naval Research and the Defense Department's Defense Advanced Research Projects Agency (DARPA). The OSTP report calls for accelerated efforts to expand interconnections among agencies.

Cost effective networking

The five year agreement between NSF and NASA took effect on January 1, 1988. The agreement states that in working toward a national internetwork system, "NASA and NSF will use the most cost-effective combination available of hardware and software technologies, communications systems, and supporting user services."

Sun sponsors Connectivity Test Demonstration

In an unprecedented showing of networking capability, over 50 computer hardware and software vendors from around the world met at Sun Microsystems in Mountain View, California for a week long, heterogeneous, test demonstration starting January 11, 1988. Called *Connectathon 1988*, the event allowed vendors to test their implementations of Sun's Network File System (NFS) to ensure that it works with all other vendors' implementations.

Wide range of hardware

The breadth of intervendor connectivity at *Connectathon 1988* spanned personal computers, technical workstations, super-minicomputers, and mini-supercomputers -- all sharing computer resources and files transparently via an Ethernet local area network. (See diagram below). In addition to testing NFS, the participating vendors tested network-wide database access between PCs and multi-user systems, using Network Innovations' application level software link. The ability to link Macintoshes, PCs and Sun workstations was demonstrated using TOPS/NFS gateway software developed by TOPS, a subsidiary of Sun Microsystems, Inc.

NFS

NFS, the best known networking service of Sun's Open Network Computing (ONC) software, is the de-facto industry standard for transparent remote file access. More than 130 organizations, including IBM, DEC, Data General, H-P, NEC, Toshiba, Siemens, Apple Computer, and leading software firms including Lachman Associates, have adopted NFS and the ONC protocol platform.

Diskless NFS

The *Connectathon* demonstration also tested an important new service, Diskless NFS. With Diskless NFS, users can operate diskless workstations on a network of other vendors' servers, including servers with different architectures and operating systems. The initial implementation being tested at *Connectahon 1988* involves Sun workstations and servers from Sun and other manufacturers.

"Because of the popularity of ONC for network computing and sharing computer resources, it's crucial that each ONC licensee has the opportunity to test their NFS implementations with other vendors'. To that end, this year's *Connectathon* provides the widest variety of NFS implementations ever assembled under one roof," said Bill Keating, technology product manager for Sun Microsystems.

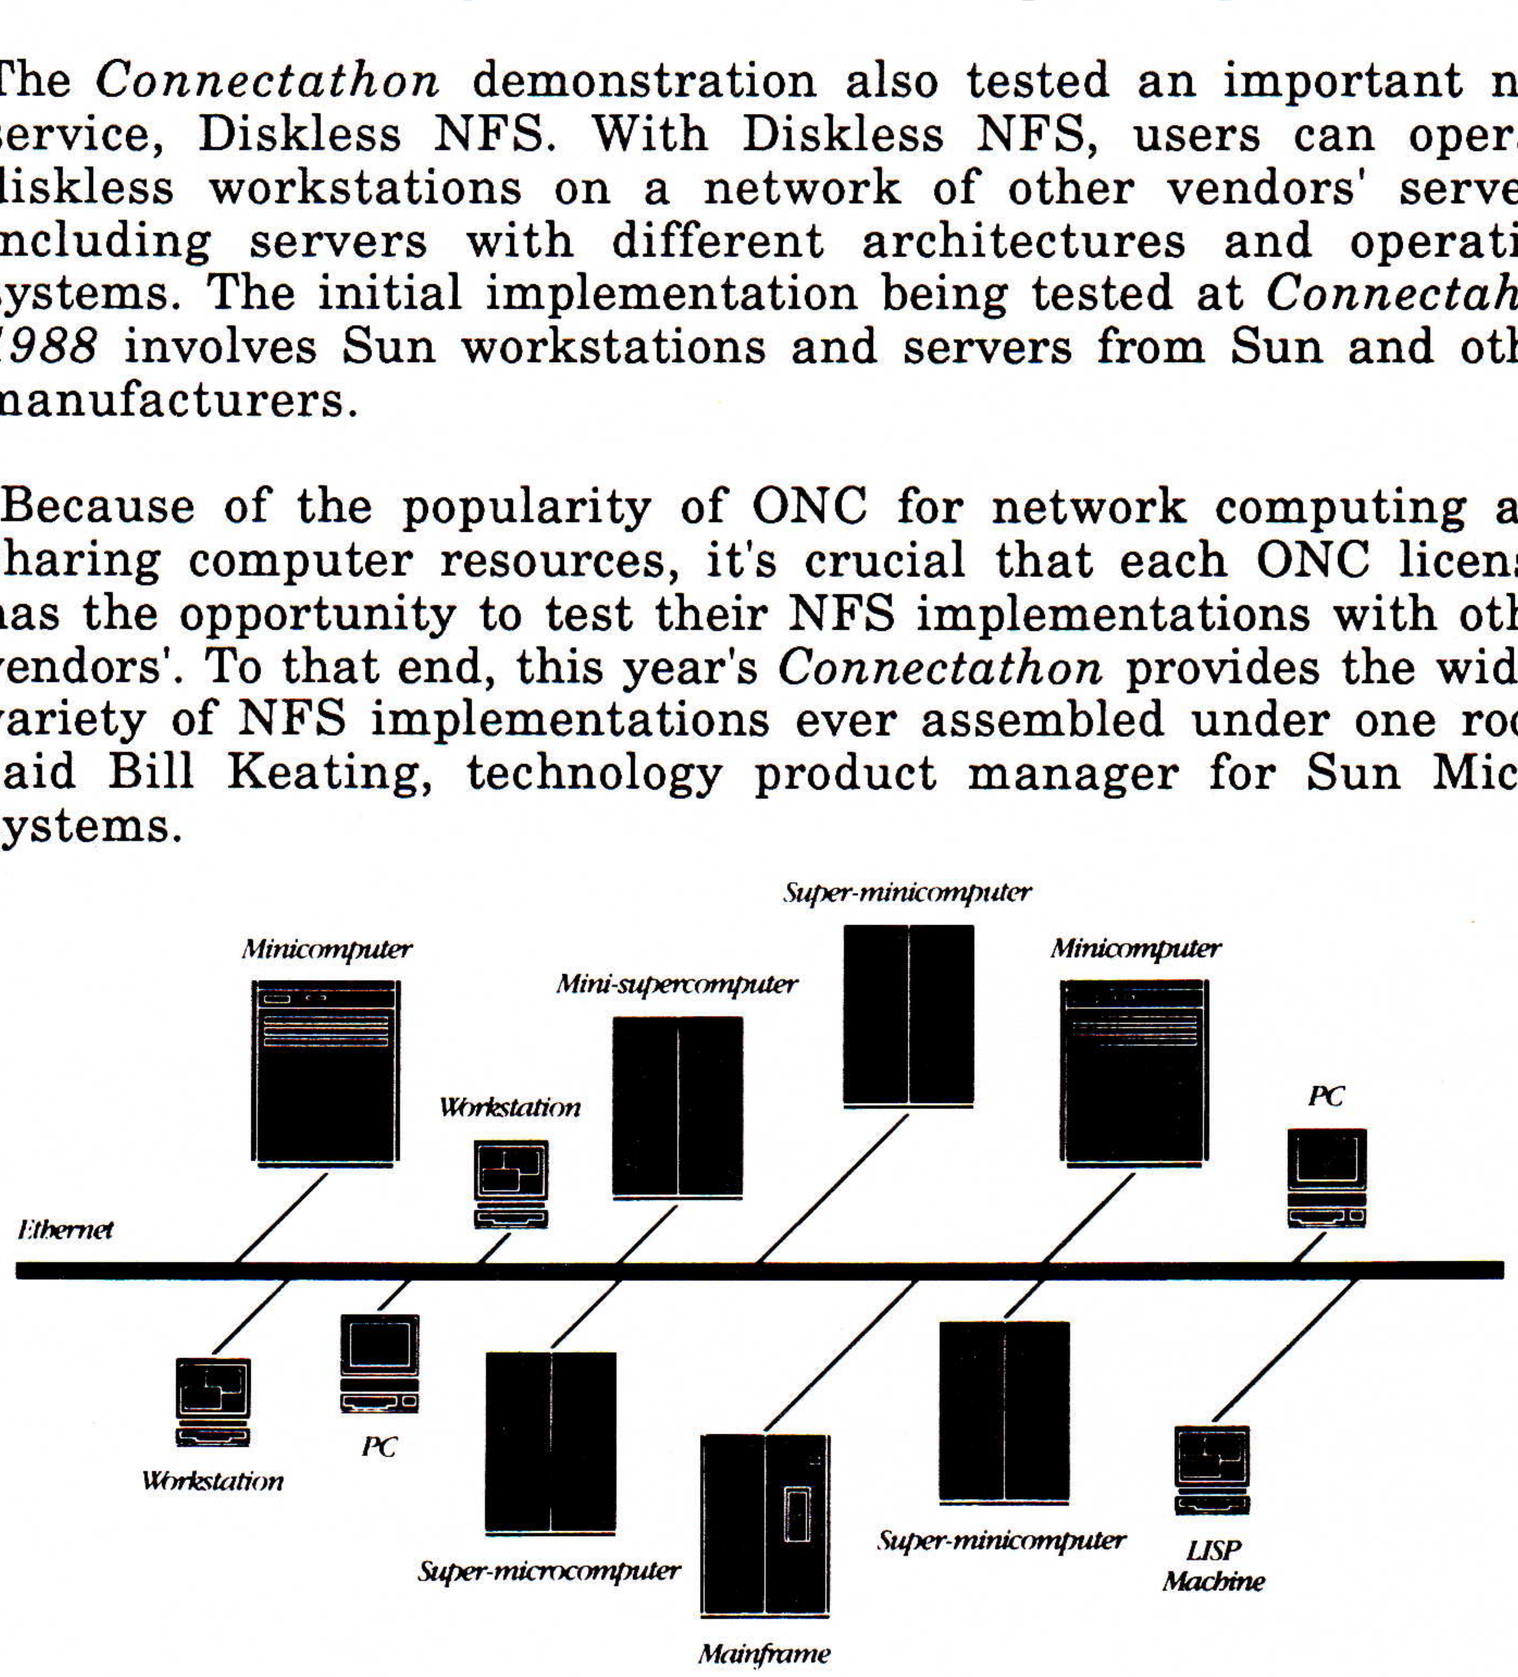

CONNEXIONS
480 San Antonio Road
Suite 100
Mountain View, CA 94040

FIRST CLASS MAIL
U.S. POSTAGE
PAID
SAN JOSE, CA
PERMIT NO. 1

CONNEXIONS

PUBLISHER Daniel C. Lynch

EDITOR Ole J. Jacobsen